Disaster = Infrastructure + Hazard

Mark Hereld,[1,2] Kibaek Kim[1,2]
[1]Mathematics and Computer Science Division
[2]Computation Institute
Argonne National Laboratory, 9700 S. Cass Ave., Lemont, IL 60439, USA
{hereld, kimk} at anl.gov

***Abstract*—We describe a framework for relating available data on hazards to impact on infrastructure. It is designed to be used by simulation platforms that may focus on a single infrastructure layer or aim to study interactions between interdependent layers. The Hazard Impact Framework (HIF) provides a flexible scheme for capturing and maintaining data on the response of elements of critical infrastructures to natural and man-made hazards. HIF also provides interfaces that enable system-wide assessment of the impact that these hazards have on the assets that construct and connect national critical infrastructures. A key use of HIF involves providing initial configurations of systems under study that include primary damage assessment across the elements of that system. These configurations can then be fed to simulation platforms to determine the level of service remaining in the damaged system, cascading effects that follow from layer interdependencies, or dynamic effects in the compromised system and to provide stochastic evaluation of likely outcomes to an event.**

I. INTRODUCTION

Understanding the impact of natural and malicious damage to critical infrastructures is fundamental to many activities important to the safety and prosperity of the nation. These activities include emergency response, development of resilient infrastructure, failure prediction, and infrastructure restoration in the wake of a disaster. Much research and development has targeted modeling and simulation of the function under strain in infrastructures important to the smooth operation of our economy and the lives of our citizenry. Moreover, attention is increasingly being paid to modeling the interactions between disparate infrastructural layers (e.g., [1], [5]). Understanding the effect of these interactions is important to gaining a systems view of the health, resilience, and vulnerabilities of critical infrastructures taken together, and in particular to understanding the cascading effects of damage to one subsystem on the entire system.

Thus far, the related research and test cases have tended to be fueled on a case-by-case basis with data describing the systems involved and the environmental influences that often drive the scenario under study. This approach yields hard-won demonstrations of principles. But until the large, heterogeneous, and incomplete datasets required become more systematically accessible, developing generally useful software platforms for routine use in real settings will be difficult.

In particular, while simulation has received much-needed and well-justified attention, less work has been done to provide systematic access to data on historical and unfolding hazards to drive the rapid configuration and execution of these simulations. Still less has been invested in providing quantitative and computationally accessible measures for the response of system components to hazards. For this reason we have developed a framework, called the Hazard Impact Framework (HIF), that provides uniform means for informing simulations of electrical, gas, water, and other key service infrastructures with an impact assessment in connection with a wide range of hazards that inflict damage on them.

The development work described in this paper is part of the Resilient Infrastructure Initiative at Argonne National Laboratory [2], [6], a comprehensive program to enable the resilient design of future infrastructure systems.

II. DATA-DRIVEN PERSPECTIVE

HIF supports simulations of one or more interacting infrastructural systems with one or more potentially damaging hazards by providing the means to couple the action of hazards on the components and function of infrastructures. The description of these hazards and the infrastructural systems they may threaten is one source of data that HIF draws on to provide its damage assessment to downstream simulation clients. The central data component of HIF, however, is a curated collection of minimal asset descriptions and linked analysis methods to assess the fragility of each asset to a given hazard. By *curated*, we mean a database built over time to include these elements so that they may be reused profitably by many applications for many simulations.

We have come to this particular point of view and design choice by noting that this connection between asset and fragility is difficult to create from disparate and incomplete data sources and so should be done only once. In taking this approach, we hope to reduce a major barrier to analysis of complex problems of great importance. In what follows we discuss the data and design problems in terms of **assets**, **fragility**, **infrastructures**, and **hazards**. See Figure 1.

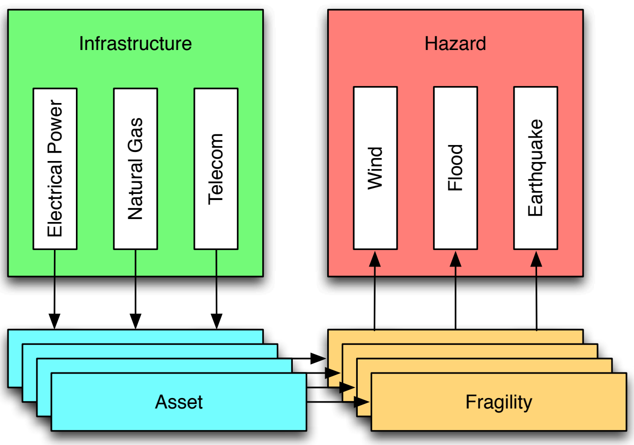

Fig. 1. The perspective taken in this paper revolves around the relationship between assets and hazards.

A. *ASSETS AND FRAGILITY*

We take as our primary perspective on the design of HIF the problem of quantitatively determining the impact of natural and unnatural forces on assets that make up the critical infrastructures under study. In practical terms this means that our problem is how to match named asset to a method for determining how much it might be damaged by particular hazard or adverse condition being considered.

An asset might be a building that houses a generator, a tower that carries high-tension lines, a section of pipeline, or a wind turbine. Given high wind conditions that prevail during a hurricane, we might want the probability that the wind turbine is incapacitated, perhaps catastrophically. With flood levels from regional rain storms rising, we might need to know whether a generator will succumb to the water.

Data describing buildings, structures, equipment, and other traditional assets of a critical infrastructure such as the electric power grid come from a wide variety of sources and make up a mixed-quality, heterogeneous, and incomplete description of the properties needed to reliably assess damage that might be inflicted on them by these hazards and threats. Moreover, some of the data is proprietary, making the data problem difficult to manage, a situation that can impede development of useful software applications.

Several sources of data, such as Federal Emergency Management Agency's methodology for estimating losses as captured in Hazus software can be used here. Summary data for different classes of structures, mostly buildings, are available in the form of fragility curves that enable probabilistic estimates of damage for various conditions. Other specifications of fragility might be stated as a simple threshold: an asset might be considered nonfunctional if the water level rises above 4 feet.

The asset descriptions are provided in a logically separate data resource that includes information about what kind of asset it is and its function and physical characteristics. This object description provides the semantic means for linking the object to its quantitative response to the conditions presented by the hazards.

The asset fragility data provides a quantitative mapping between the local effect of the hazard to the viability of the asset. For example, wind speeds in excess of 120 mph may result in a 50% chance that a building will suffer extensive damage. These data come in many forms depending on available information and common practice. In some cases the fragility is captured in a set of curves describing probability of damage at various levels of severity, while in others a simple threshold may be sufficient (or all that is available).

We note that focusing on assets and fragility descriptions will enable us to consider a wide range of systems. Assets are not limited to the physical objects represented by the nodes and edges of network descriptions of, say, the electrical grid. As long as an asset can be paired with a behavioral response to a hazard or threat, it can be handled by HIF. This generality opens the way to considering capabilities, capacities, commodities, and other less tangible components that might be included in models. For example, an agent-based model might include "access" to a roof as an asset that is important to the completion of important tasks.

B. *INFRASTRUCTURE AND HAZARDS*

As noted, in the perspective taken here an infrastructure is modeled as a collection of assets with behaviors that are pertinent to the functioning of the infrastructure and are the subject of simulations. The infrastructure layer often provides a description of the geographical layout of the principal components of a critical infrastructure, such as the electric power grid. In this case, a shapefile might provide the locations of generators, substations, towers, and other components, as well as the connectivity between them provided by the power lines.

Often one wishes to represent a particular modeled infrastructure as a network with nodes and edges, each of which may represent assets subject to damage by the action of the hazard. Although several infrastructural systems of interest can be modeled in this way — including the electrical power grid, the natural gas distribution network, transportation and waterway networks, and telecommunications — this approach is not necessary for HIF as long as the model can be cast in terms of assets and fragility measures.

We note that much of the data describing these details of critical infrastructure are proprietary, not widely available, and not uniformly specific and detailed. Naming inconsistencies, for example, present problems for robust identification of assets across data sources.

Examples of hazards to be considered include earthquakes, high wind, and flooding. Available data about historical and ongoing hazards of these sorts are maintained by various government agencies and often packaged for relatively easy access in standard formats from the web. For example, a hurricane description might include a shapefile giving the wind strength over an affected geographical area.

As noted in the discussion of fragility assessment, the form these data take will determine the appropriate analyses. For example, wind speed data and earthquake displacement data are often available as a two-dimensional scalar field implicitly represented by isolines in a shapefile. Flood data, on the other hand, can be specified in terms of probabilities that water will achieve a stated depth. Each representation offers different opportunities for providing damage assessment.

C. DESIGN GOALS

The design of HIF has been driven by several goals. Our first and highest-level goal has been to automate the assessment and analysis process for rapid evaluation of hazard impact on coupled infrastructural layers. Several advantages follow from achieving this goal. The automated process will result in faster time to prediction, which will in turn allow for more time to consider the best response and enable quicker overall response. Several problems need to be solved to enable this automation. Among them are development of reliable and flexible interfaces to a variety of data sources and creation of a flexible scheme for representing probabilistic descriptions of infrastructural responses to stresses induced by the hazards.

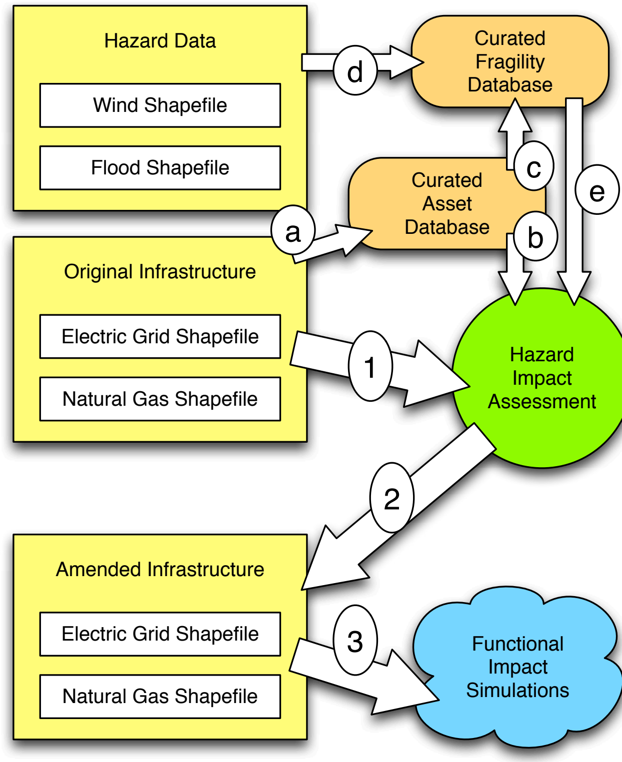

Fig. 2. Overview of the Hazard Impact Framework that performs a failure analysis used to generate the initial conditions for downstream simulation engines. At the heart of HIF is the assessment engine that is driven by data describing hazards, infrastructure layers, and curated data laying out the properties of assets and their quantitative vulnerability to hazards.

A second goal of the design and development work has been to enable flexible configuration of situations involving many moving parts. For a given hazard, any of several infrastructures may be at risk. Therefore, HIF must be able to seamlessly connect the local effect of the hazard on the environment to its impact on any of a wide range of infrastructural assets.

The third major goal of the development process has been to make the capabilities of HIF easily extended. The software design and development process needed to support straightforward addition of the necessary data and code in order to enable HIF to respond to a new or improved hazard type. It also needed to support low-cost addition of infrastructure components, descriptions of the impact of a hazard on their functionality, and even whole new infrastructural layers if required.

Additional objectives include ability to work with appropriately large geographical regions when considering the impact of a hazard on infrastructure, which will require consideration of data and computation scale; ability to handle time at various granularities in order to facilitate capture of dynamic processes, effects of

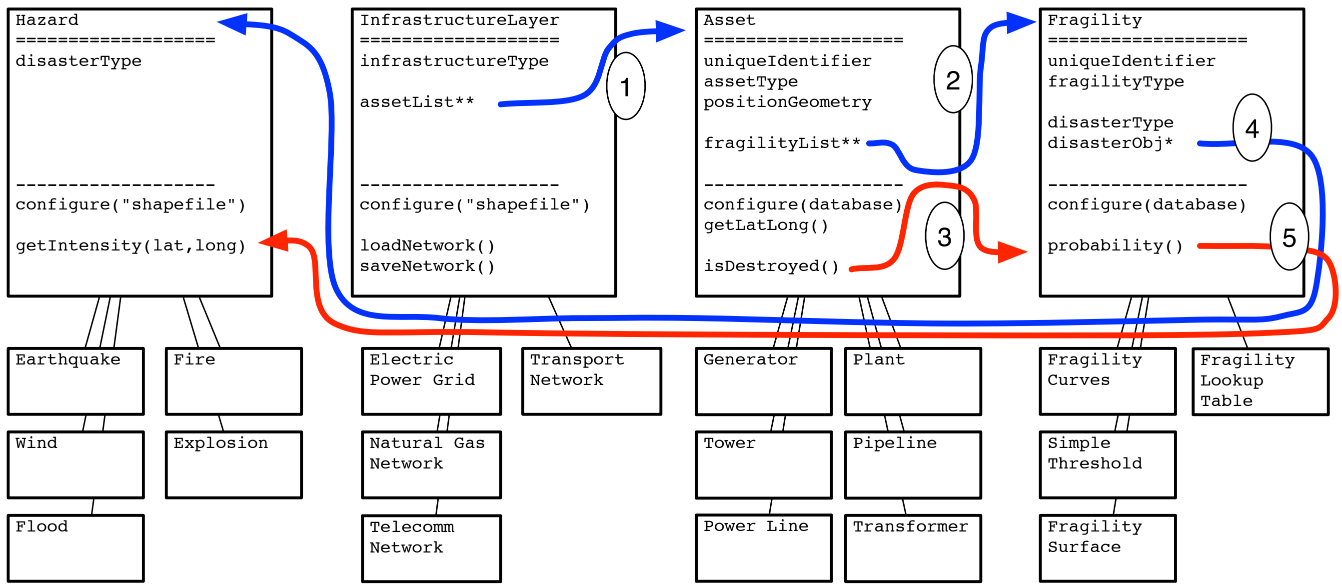

Fig. 3. Diagram showing the class structure of HIF. The four principal classes show how they support the methodical assessment of damage due to a hazard scenario. Each abstract class at the top is customized in subclasses, examples of which are found in the smaller boxes below.

mitigation, and cascade effects; and leveraging of open source technologies wherever possible in order to control cost, and foster flexibility portability.

III. HIF Process Flow

The high-level view of the failure analysis that HIF facilitates is shown in Figure 2. Its role is to assess the first-order effect of the hazard(s) on the infrastructural layers being considered, which can then be used by simulation applications. The undamaged infrastructure configurations initialize the HIF process at (1) in the figure. Its chief output at (2) is a description of the anticipated initial state of the infrastructure as impacted by the hazard – which lines are down, which plants are damaged and no longer online, and so on. This description of the diminished infrastructure can be passed at (3) to dynamical simulations for detailed study of the consequent coupled behavior of the system.

The hazard impact assessment, which leads directly to a failure analysis, is driven by the files that describe the four principal aspects of the scenario: hazards, infrastructure layers, asset descriptions, and asset fragility data. Assets found in the infrastructure description files are matched at (a) in Figure 2 to entries in the curated asset database. These are passed at (b) to the assessment process and are used at (c) to identify the appropriate fragility measure for each hazard that may have impact on the asset. The hazard data at (d) is used to evaluate the resulting damage assessment, which is passed at (e) to the orchestrating processes. After this process has been carried out on all the assets under consideration, the resulting amended infrastructure configuration is made available at (2).

Note that this organization of the data means that the detailed description of an infrastructure is actually split into three separate data perspectives, roughly: layout or organization, component assets, and fragility of assets. The data are purposefully separated along these functional lines to facilitate curation, reuse, and flexible creation of new scenarios. These categories also facilitate development of a well-organized software architecture that supports rapid incorporation of new hazards, infrastructural components, and fragility descriptions.

IV. Implementation

Figure 3 outlines the principal ingredients of the software framework. The `Hazard` class gives access to data and analysis specific to each kind of hazard that might be included in the framework (hurricane winds, earthquake, flood, etc.) through a suitably generic interface to enable hazard-agnostic code to assess its impact on the infrastructural components. In this way, the failure analysis can be applied transparently across hazard types, asset types, and the many ways that fragility is described in the discipline.

A scenario consists of one or more hazards, each instantiated as a `Hazard` object. The scenario also generates one or more `InfrastructureLayer` objects and fills it with data from input shapefiles. These layers maintain a list of component `Asset` objects, illustrated by the blue line (1) connecting the `assetList` data structure to each instance of `Asset`. Likewise, each `Asset` object maintains a list of response functions, one for each type of hazard in the scenario. These `Fragility` objects allow the asset to assess its quantitative response to the corresponding hazard. Another blue line (2) in the figure illustrates that relationship. In terms of the data structures that support the framework, each `Fragility` object knows which `Hazard` object it uses to relate the intensity of the hazard to the damage inflicted on the asset, again indicated by a blue line (4).

Thus far, all these connections are agnostic of the type within class. That is, an `InfrastructureLayer` can have many types of asset, just as an asset can utilize many different types of fragility. Also supporting this agnosticism is the design choice to encapsulate the response of an asset to a hazard in terms of three generic functions: `isDestroyed()`, `probability()`, and `getIntensity()`. The red lines in the figure indicate how these functions connect the data describing the hazards in the scenario to the failure assessment of each asset. An `Asset` object may be queried through the `isDestroyed()` function, which finds the responsible `Fragility` object and asks it for an assessment via `probability()` (3), which uses `getIntensity()` against the appropriate `Hazard` object. With this modular design and careful separation of data concerns from assessment logic, HIF can be configured to assess a wide range of scenarios by changing data, not code. Furthermore, the process can be extended to include new hazards, infrastructures, assets, and fragility descriptions.

Figure 4 shows a fictitious use case of our prototype implementation for the IEEE 64-bus system placed on the Gulf coast in the path of Hurricane Ivan. This test of HIF required two pieces of configuration data. The shapefile for Hurricane Ivan contained wind-speed isolines. The shapefile for the IEEE 64-bus reference system included buses and the network connectivity. In addition, we generated a simple generic database describing the assets putatively represented by the buses and the corresponding fragility methods to be applied. In practice, this database would be a curated resource mapping real assets and their fragilities.

HIF constructed an internal model following the procedure described in Section III. Assessing the impact on each bus caused one to be knocked out and marked with a red `x` if the wind speed at the location exceeded the asset's fragility criterion. The green and orange lines in the figure represent the locations with the wind speed of 74 and 58 knots, respectively. The evaluation scheme provided by the `probability()` method of the assigned `Fragility` object performed appropriate interpolation or extrapolation of these isolines to the bus position as needed.

In a more elaborate test, we looked at the impact of flood water level and hurricane wind strength on the electric power grid and the natural gas infrastructure of a small geographic region. After HIF generated an assessment of the net damage to the assets involved, the resulting description of the reduced capacity system was sent to an iterative solver using EPfast [4] and NGfast [3] alternately to identify additional asset knockouts resulting from dependencies between the damaged electric and natural gas systems [5].

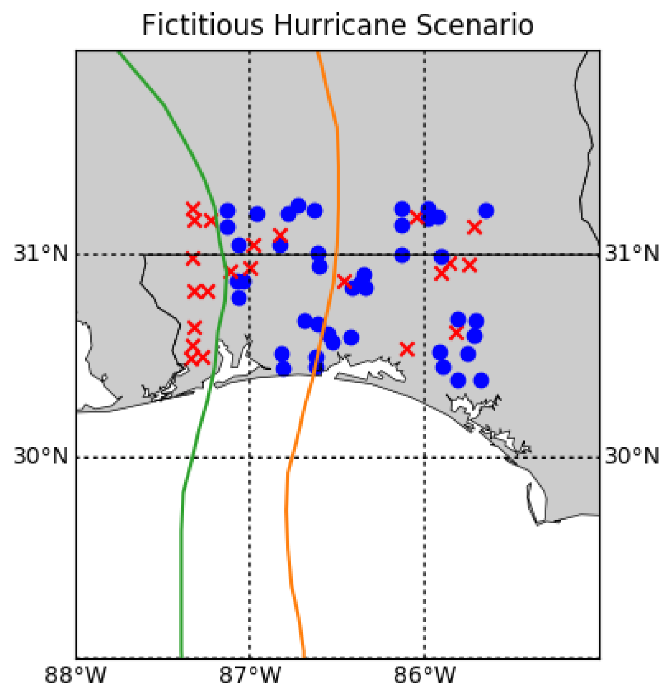

Fig. 4. A fictitious scenario in which the IEEE 64-bus system data are placed on the Gulf coast in the path of Hurricane Ivan. Each knocked-out bus is marked with a red ×.

V. DISCUSSION AND CONCLUSIONS

The new process enabled by HIF begins with configuration of a scenario. The configuration, contained in a text file, includes the names of data sources required to describe the hazards, infrastructure layout, properties of the individual elements of the infrastructure, and the quantitative response of these to any hazard that is part of the scenario. These data sources are in standard and widely used formats where possible (chiefly the ubiquitous shapefile format), and in simple human-readable text files where no suitable standard exists.

Agnosticism through object-oriented design in HIF leads to clear demarcation between details specific to a given asset's response to a particular hazard and the general notion of whether and by how much it has been degraded. This agnosticism is achieved by exchanging generic quantities between the various software components where possible, for example through geophysical location in latitude and longitude, geometrical description of geophysical extent, and probability that an asset has been damaged beyond usefulness. This approach is in contrast to passing physical quantities that are specific to a hazard (wind, displacement, surge level) or asset (tensile strength of power lines, height of electrical boxes, structural strength of roof and walls). The ultimate connection between the physical units is done in this framework at the bottom of the layered code.

In addition to this simple test case, HIF enables a number of capabilities that would otherwise be too impractical to be exercised routinely. It is now possible to do the following.

- Rapidly test the impact of hypothetical scenarios by substituting data sources at configuration time,
- Establish detailed quantification of the uncertainties attendant to any predicted impact by running ensembles of HIF configured simulations to propagate the probabilistic damage assessment provided by HIF into the consequent range of functional outcomes determined by the simulation over distinct initial configurations,
- Study ensembles of scenarios,
- Execute searches through infrastructure configurations to probe potential weaknesses,
- Provide analyses to guide infrastructure planning strategies.

By capturing hard-to-get data and methods describing assets and their behaviors under strain and packaging these into easy-to-access forms, the Hazard Impact Framework makes a wide range of powerful applications practical and easy to implement.

ACKNOWLEDGMENT

We acknowledge many useful discussions with Brian Craig, Edgar Portante, Stephen Folga, Edwin Campos, Megan Clifford, Guenter Conzulman, Diane Graziano, Charles Macal, Jonathan Ozik, Frederic Petit, Julia Phillips, Angeli Tompkins, and Victor Zavala. This material is based upon work supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357.

REFERENCES

[1] Chiang N.Y., and Zavala, V.M. "Large-Scale Optimal Control of Interconnected Natural Gas and Electrical Transmission Systems," Applied Energy 168, 226-235 (2016).

[2] Clifford, M. "National Call to Action: The Resilient Infrastructure Initiative," George Mason University Center for Infrastructure Protection & Homeland Security (December 2015).

[3] Portante, Edgar C., Brian A. Craig, and Stephen M. Folga. "NGfast: A Simulation Model for Rapid Assessment of Impacts of Natural Gas Pipeline Breaks and Flow Reductions at US State Borders and Import Points." In Proceedings of the 39th conference on Winter simulation: 40 years! The best is yet to come, pp. 1118-1126. IEEE Press (2007).

[4] Portante, Edgar C., Brian A. Craig, Leah Talaber Malone, James Kavicky, Stephen F. Folga, and Stewart Cedres. "Epfast: A Model for Simulating Uncontrolled Islanding in Large Power Systems." In Proceedings of the Winter Simulation Conference, pp. 1763-1774 (2011).

[5] Portante, Edgar, Brian Craig, Jim Kavicky, Leah Talaber, and Stephen Folga. "Modeling Electric Power and Natural Gas Systems Interdependencies." The CIP Report (2016).

[6] https://www.anl.gov/egs/group/resilient-infrastructure, Resilient Infrastructure Initiative, Argonne National Laboratory (URL visited May 1, 2017).